The way of AI: paths, pitfalls & tools
in surgery, medicine and health care

André Carrington PhD PEng

December 5, 2019

The Ottawa Hospital RESEARCH INSTITUTE | L'Hôpital d'Ottawa INSTITUT DE RECHERCHE

Affiliated with • Affilié à uOttawa

Statistics, ML and AI

Thinking rationally

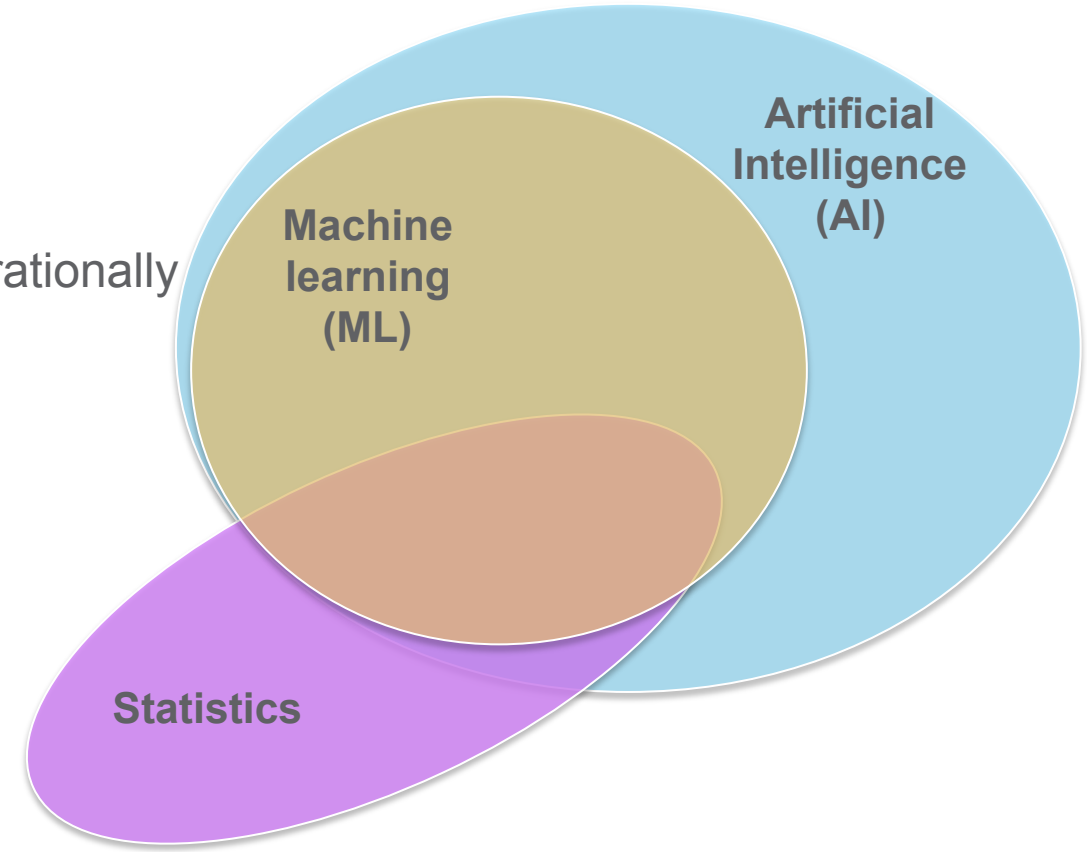

Thinking rationally or humanly; or

Acting rationally or humanly or autonomously.

The Ottawa Hospital RESEARCH INSTITUTE | L'Hôpital d'Ottawa INSTITUT DE RECHERCHE

Thinking/acting, humanly/rationally adapted from Russell and Norvig

Statistics, ML and AI

Single task (mostly)

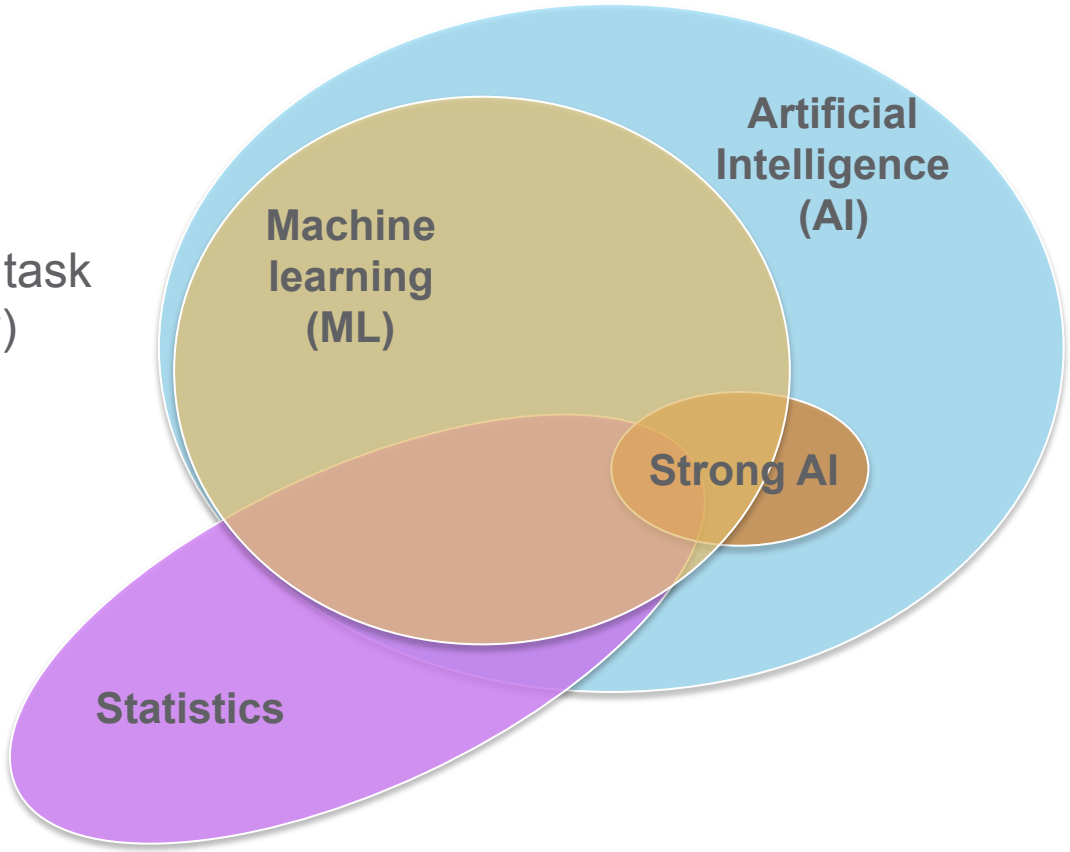

Multi-task extreme

Artificial general intelligence (AGI)

Statistics, ML and AI

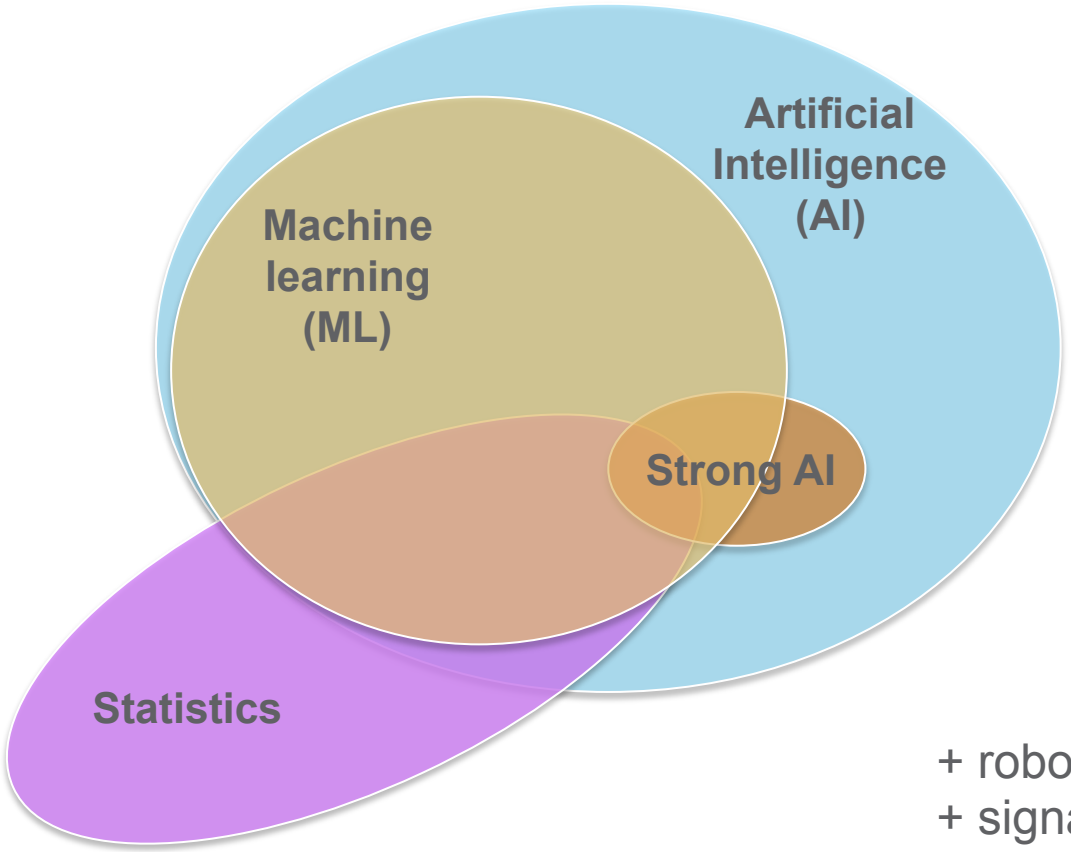

+ robotics (control theory)
+ signal processing
+ expert systems

Other overlapping concepts/terms

- Deep learning
- Data science
- Data mining
- Analytics
- Clinical decision support systems
- Image processing
- Computer vision
- Computer aided detection (CAD) or diagnosis (CADx)
- Business intelligence
- Statistical learning
- Pattern recognition
- Operations research
- Knowledge based systems


One day Alice came
to a fork in the road
Statistics
AI


Statistics
AI


Cat: where do you want to go?
Statistics
AI


Cat: where do you want to go?
Alice: to paradise for analysts.


Cat: where do you want to go?
Alice: to paradise for analysts.
do both paths go there?

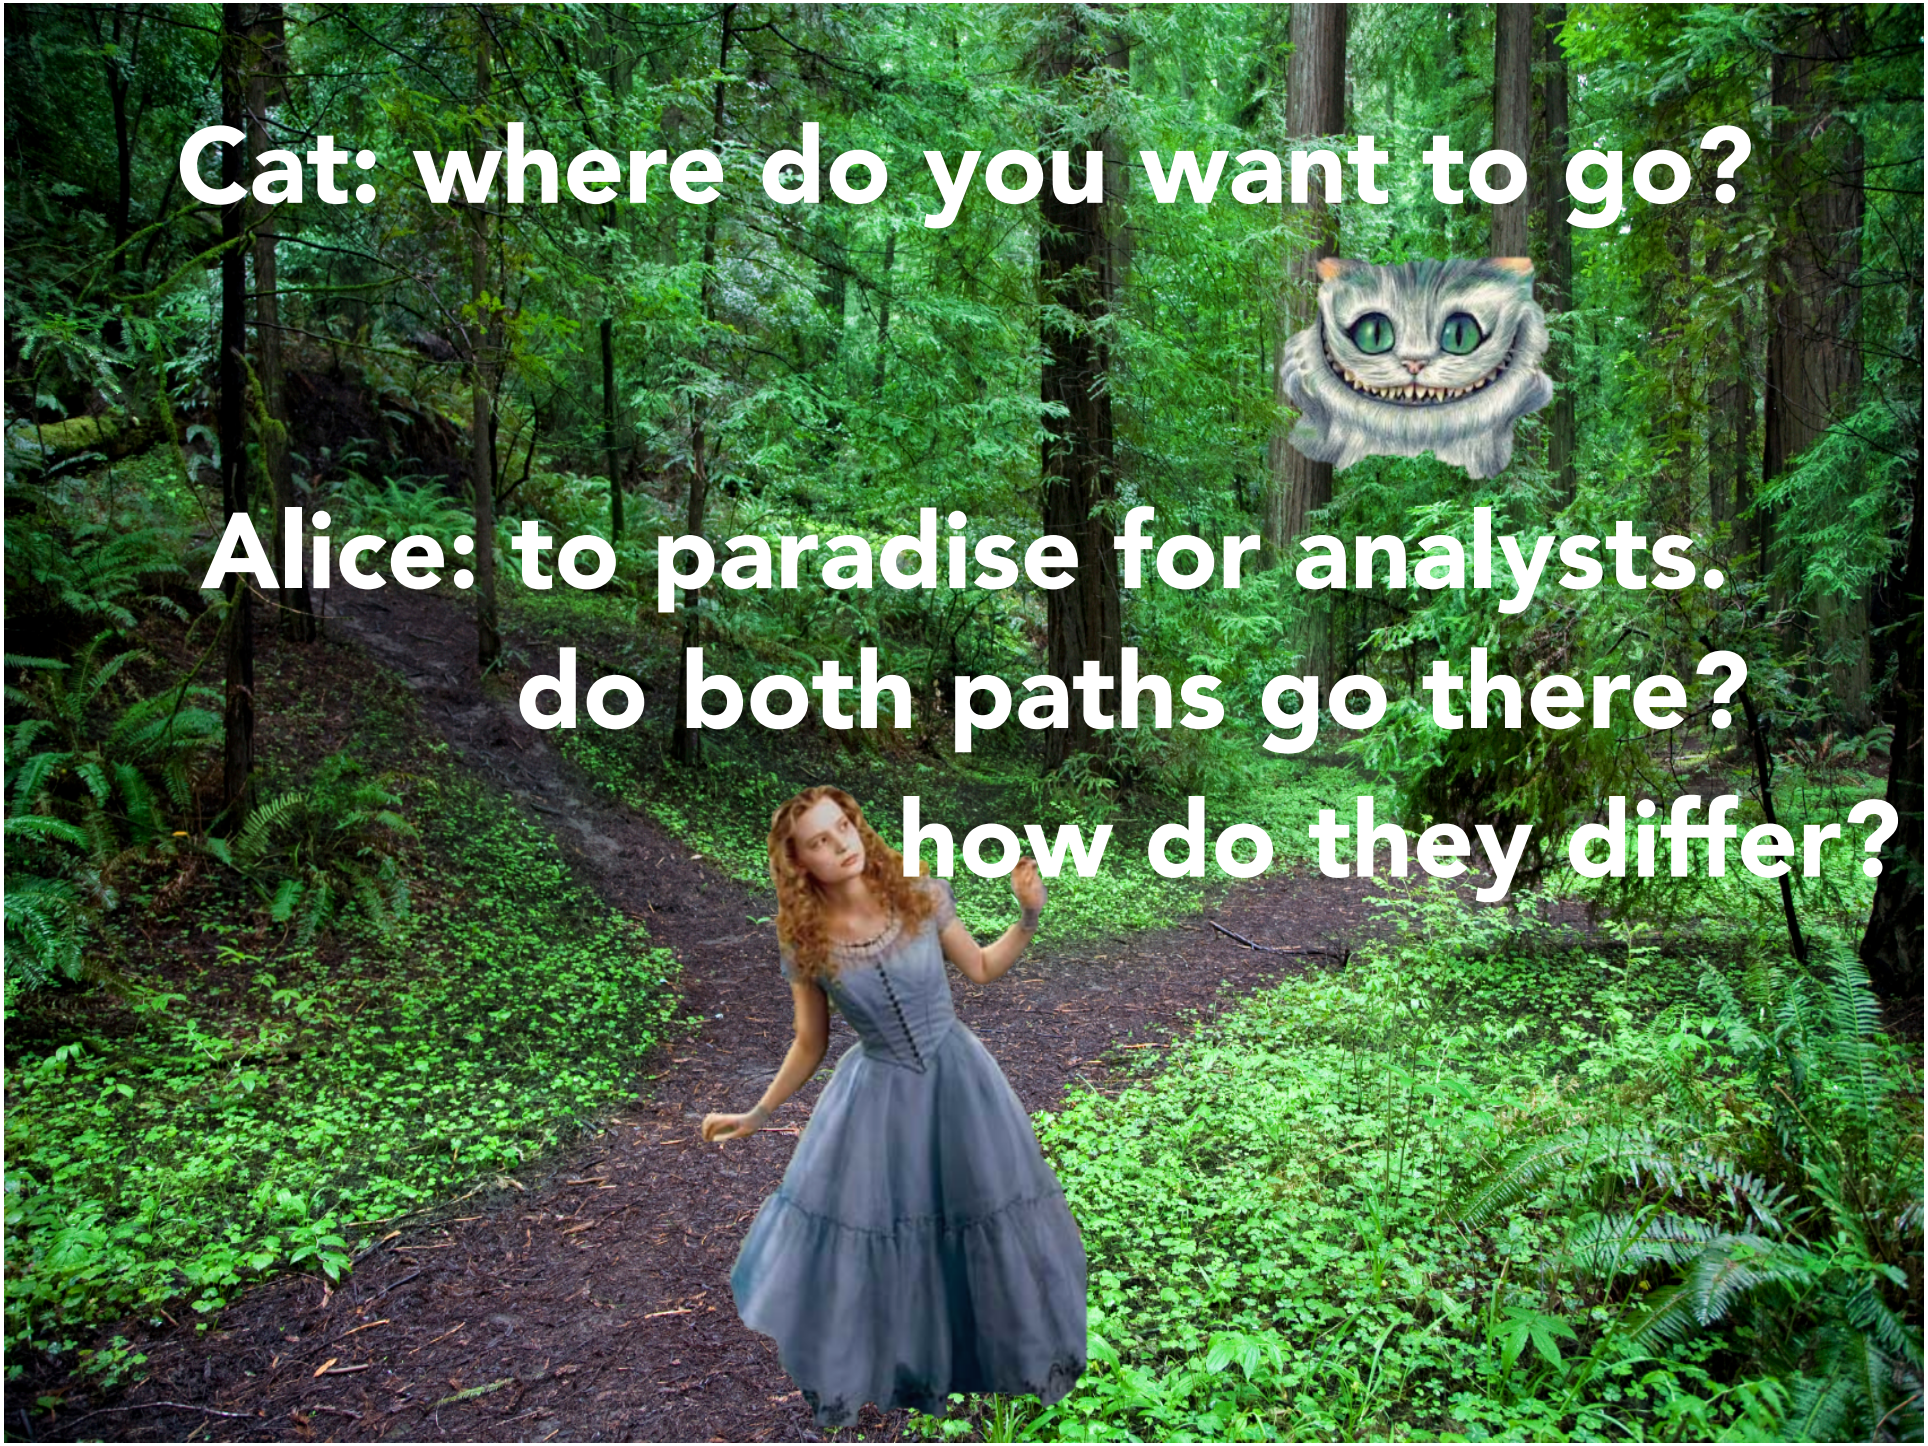
Cat: where do you want to go?
Alice: to paradise for analysts.
do both paths go there?
how do they differ?

What is paradise (for a task)?

- Automation
- Discrimination
- Calibration + calibrated scores
- Statistical validity

What is paradise (for a task)?

- Automation
- Discrimination
- Calibration + calibrated scores
- Statistical validity
- Transparency
- Explainability + causal insight = causability
- Usable/available tool
- ➢ **Clinical utility**

Do Statistics and AI paths both go to paradise?

YES, for some tasks

- Classification
- Regression
- Dimension reduction
- Time series
- Clustering
- Anomaly detection
- Sequential decision-making

Do Statistics and AI paths both go to paradise?

YES, for some tasks

- Classification
- Regression
- Dimension reduction
- Time series
- Clustering
- Anomaly detection
- Sequential decision-making

MAYBE? Survival analysis

Stats and machine learning can use

- Numeric, binary, categorical, ordinal, censored, ratios

Age	Sex	Race	Cancer stage	Prognosis
25	M	Black	II	4

- Numeric waveforms, spectra, ICD-10-CA, HL7

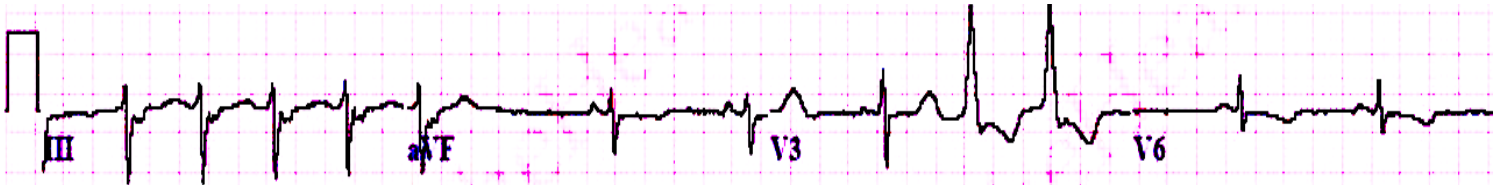

…extracting features

Machine learning can use

- **2D/3D images as-is** (no extraction)
- **genomic sequences**
- **natural language (freetext)**
- robotic path planning
- graphs or networks
- speech recognition

Do Statistics and AI paths both go to paradise?

Statistics avoids:

- Image processing*
- Biologic sequence analysis
- Natural language processing
- Speech recognition
- Speech generation
- Graph theory
- Game theory

Do Statistics and AI paths both go to paradise?

Statistics avoids:

- Image processing*
- Biologic sequence analysis
- Natural language pro
- Speech recognition
- Speech generation
- Graph theory
- Game theory

AI avoids:

- Imputation
- De-identification
- Kernel density est.
- Descriptive statistics
- Performance measures

Do Statistics and AI paths both go to paradise?

Statistics:

- Needs more analysis of data/model
- More health focus
- Needs less data

AI:

- Better than stats for very complex problems
- More mindshare
- Needs more data

Do Statistics and AI paths both go to paradise?

Statistics:

- Needs more analysis of data/model
- More health focus
- Needs less data (model driven) global model
- Usually easier to explain, inspect

AI:

- Better than stats for very complex problems
- More mindshare
- Needs more data (data driven) local model
- More effort, expertise to explain, inspect

Data in biostatistics

- Structured data
- Big or small data
- Thin (few features, many samples), not wide
- Static or longitudinal
- Feature collinearity avoided
- Feature parsimony sought

Data in machine learning

- Structured **or unstructured** data
- Big or small data
- Thin **and wide**
- Static, longitudinal, **or online as an incoming stream**
- Feature collinearity **okay / better tolerated**
- Feature parsimony **not needed: thousands allowed**
- Calibration (goodness-of-fit) **often neglected**

Learning in biostatistics

- Maximizes likelihood
- Calibration (goodness-of-fit) important
- Converted to a single data type
- A few mixed data methods
- Can explicitly include noise & error variables (uncertainty)
- Exceptions: LASSO, Ridge

Learning in machine learning

- Maximizes **posterior probability** and may
  - **Maximize margin**
  - **Control complexity** with sparsity & shrinkage
- Goodness of fit and p-values **often ignored**
- A few more mixed data methods
- Sometimes bio-inspired

Pitfalls and challenges

- Collaboration assumptions: e.g., terms, methods
- Non-bias/cost assumptions: e.g., AUC relevance
- Data assumptions: e.g., PCA - linear
- Expert cameras and bake-offs: e.g., linear SVM
- Being unFAIR: e.g., re-usable open science

Collaborating challenge: overlapping terms

features ML, covariates S, attributes S,
independent variables S, random variables S
explanatory variables S

target ML, outcome HC,
dependent variable S, output ML
response variable S

	wgt	height	pulse	age	sex	acr	gfr	ckd	
instances ML, samples S, observations HC,S, cases HC	170	68	80	65	0	30	60	1	**events** HC,S, positives
	150	65	60	46	1	5	30	1	
	155	66	65	22	0	2	95	-1	negatives
	160	68	60	37	1	2	100	-1	

data matrix ML, feature matrix ML
sample S

Collaborating challenge: different perspectives and methods

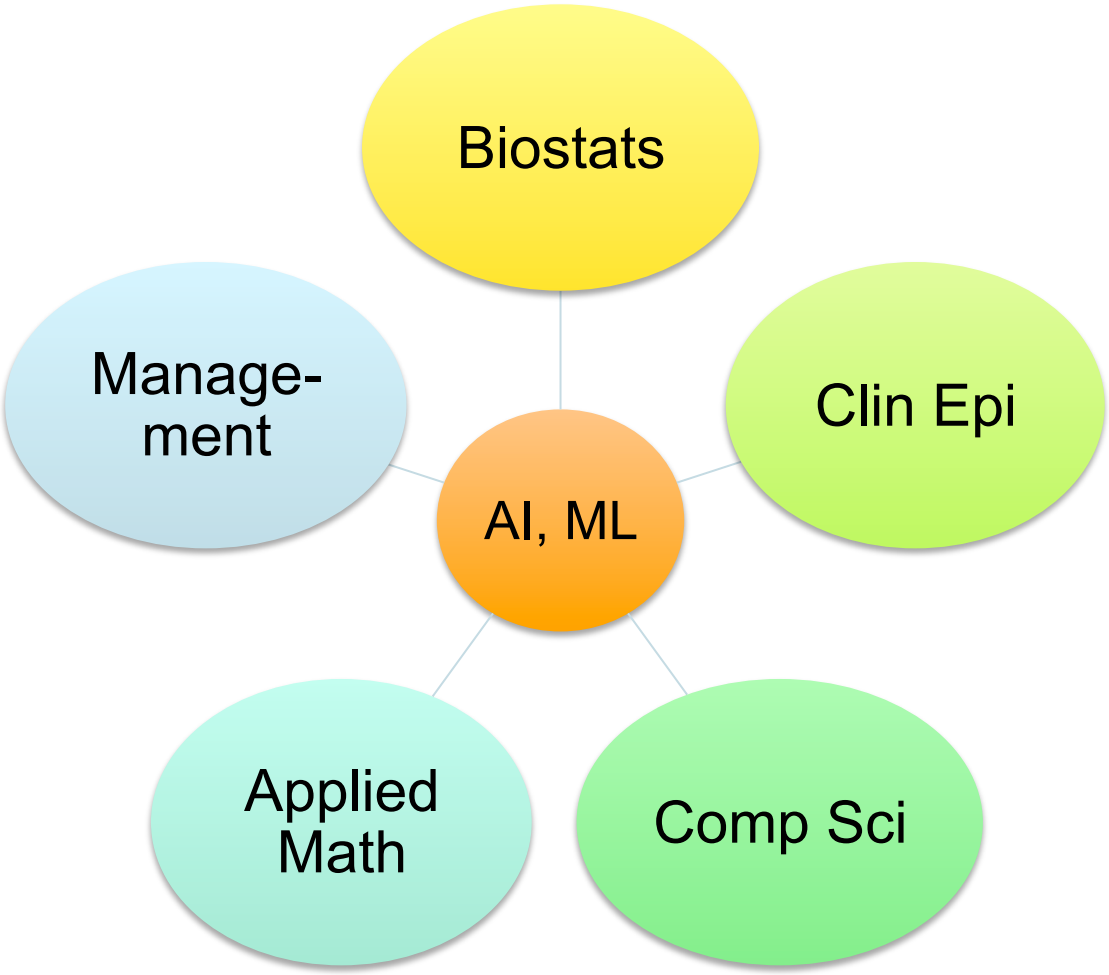

Area Under the Curve – relevant regions

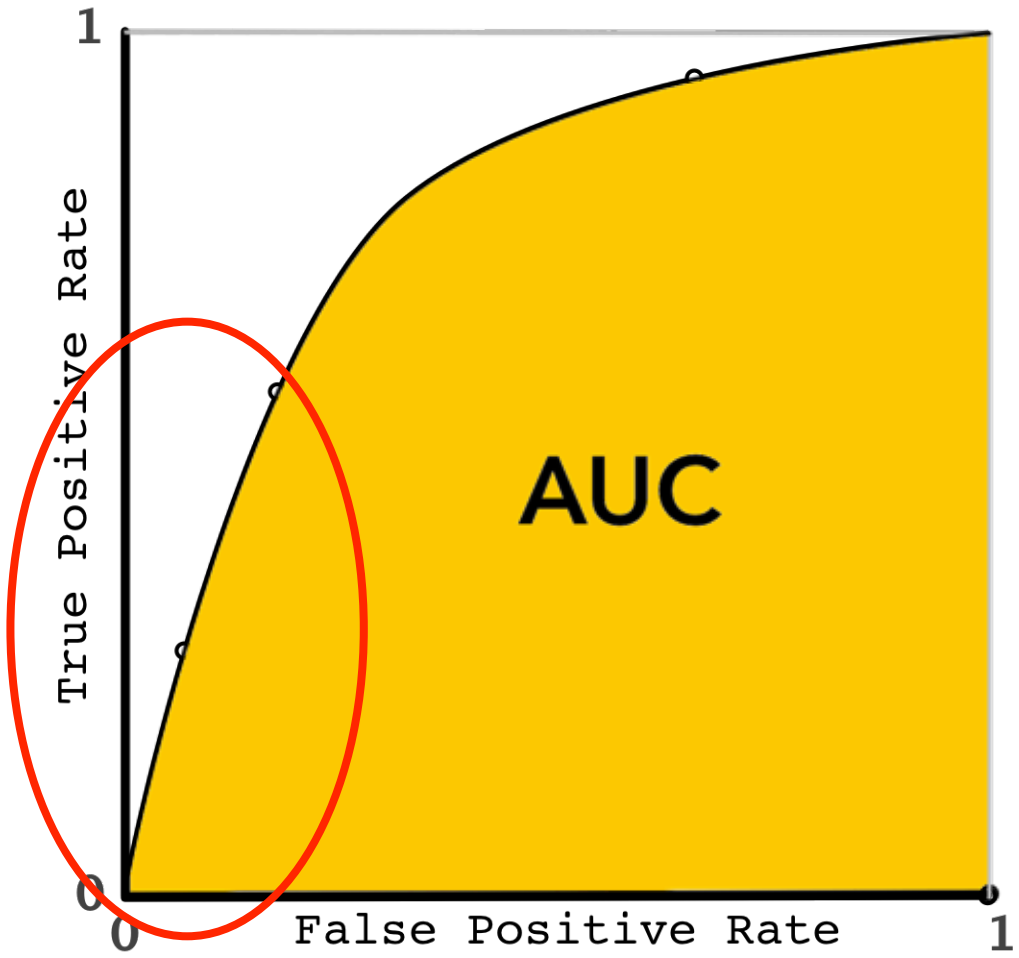

Using the area under the ROC curve (AUC)[29].

But only some regions are relevant![4,5,6]

For low prevalence, the region of interest is at left[7,8].

A pedestrian was killed by an uber robocar

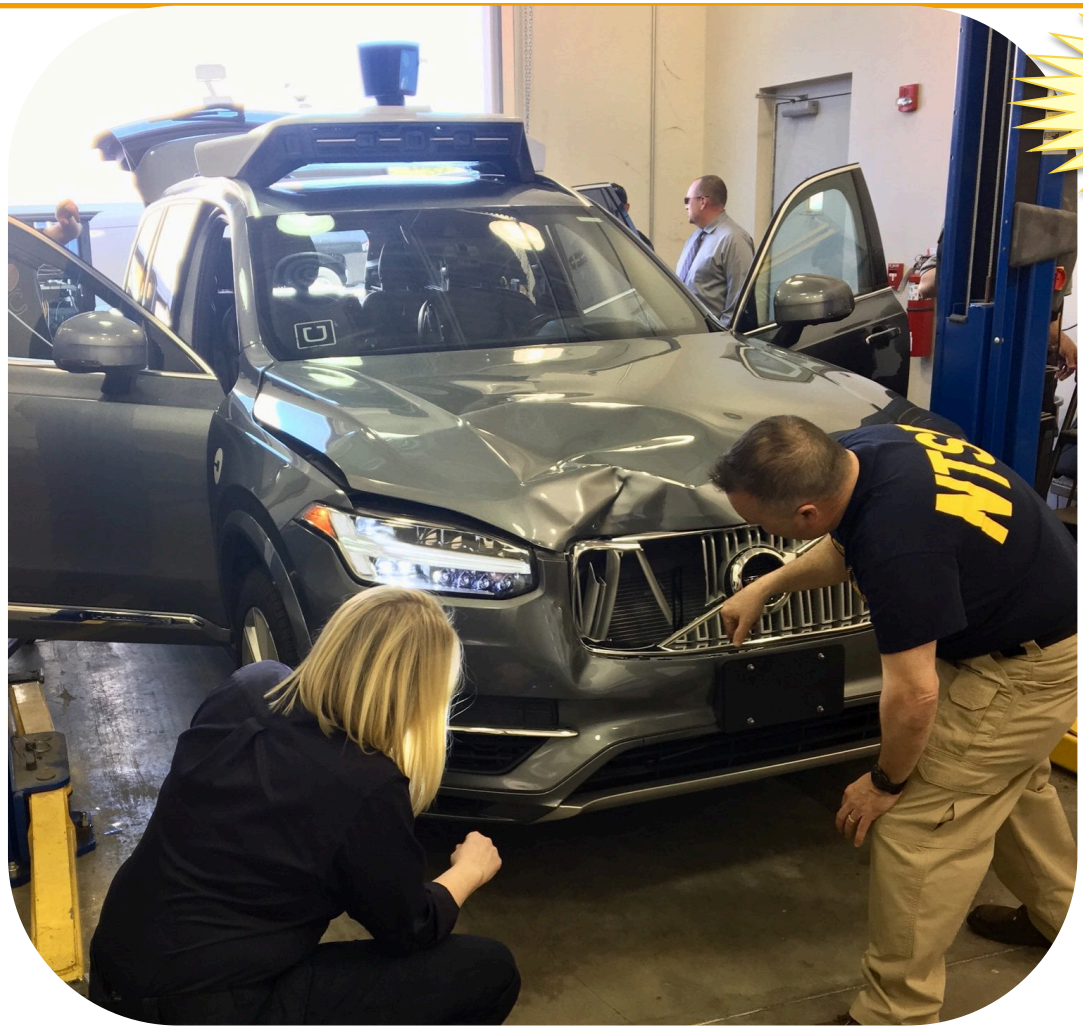

Hidden slide

Medical device security

Hidden slide

- Organizations that share experience and information on security, called ISAOs/ISACs are new in medical devices, as of 2017.
- In defense, finance and energy they have existed for decades.
- We need a margin of safety!

AGENDA

1. Paths
2. Pitfalls and challenges
3. Tools and resources

You are not alone!

- AI and ML are an explosion of concepts
- Wider than statistics and deep
- Weka has 180+ methods

You are not alone!

- AI and ML are an explosion of concepts
- Wider than statistics and deep
- Weka has 180+ methods

plan, study, try, then

learn as you go.

Foundational skills

to understand and do AI/ML well, are:

- algebra (advanced)
- statistics (advanced)
- ML methods (advanced) and pipelines
- programming, databases, data wrangling
- subject matter concepts, processes, data
- explainable and equitable AI
- some probability, info theory, calculus, optimization

Additional/optional skills

to do/use some AI/ML methods, are:

- some optimization (advanced)
- some functional analysis and topology
- some thermodynamics and neuroscience
- some causality, theories of science and evidence
- some learning theory

I can help with

- Questions, papers, books, authors, conferences, forums and mailing lists, contests, benchmark data

- **I cannot help with:**
  - online AI courses
  - courses at uOttawa
  - introductory AI books

My experience:

in-person

uWaterloo

none avail my time

Example graduate courses (at UW)

Hidden slide

- Introduction to Mathematical Oncology
- Statistical Learning - Classification
- Logistic Regression and Its Application
- Correlation and Regression
- Computational Linguistics
- Advanced Pattern Recognition
- Statistical Image Processing and Multidimensional Modeling
- Applied Machine Learning
- Numerical Algorithms and Image Processing
- Medical Image Processing
- Computational Techniques in Biological Sequence Analysis
- Health Informatics II - Application Domains
- Health Informatics I - Data Structures and Standards
- Human Aspects of Software Engineering

Example undergrad courses (at UW)

Hidden slide

- Undergraduate research assistant: Models of relative visual sensitivity in optometry (multivariate non-linear regression)

- Systems Design Workshop II: Autonomous guided vehicle with ultrasound and fuzzy logic control
- Systems Design Workshop I: Music parsing (pattern recognition)
- Analysis of Large Systems
- Linear Systems
- Numerical Analysis
- Engineering Optimization
- Software Engineering

- Systems Models
- Linear Systems and Signals
- Statistics
- Differential Equations
- Probability
- Calculus II
- Calculus I
- Linear Algebra
- and more

Papers on AI in general

Details included

- Rajkumar et al – AI in Medicine 2019 — analysis
- Upshur - AI ML…impacts on…Family Medicine 2019 — perspective
- Leo Brieman – Statistical Modeling: The Two Cultures 2001 — seminal
- Mukherjee – AI vs. MD (New Yorker Magazine) 2018 — layperson
- Topol – High performance medicine… 2019 — current
- Cruz and Wishart - Survey of ML in Cancer Prediction… 2006 — good review
- Wu et al - Top 10 Algorithms in ML 2008 — historical
- Kotsiantis - Survey of Supervised ML 2007 — useful
- Bhaskar et al – Survey of ML for Bioinformatics — DNA

Papers on AI in surgery

Details included

- Hashimoto et al - AI in surgery promise and perils 2018
- Kassahun et al - Surgical robotics… 2019
- Bernardo - VR and simulation in neurosurgical…2017
- Winkler-Schwartz et al – AI…to assess surgical expertise in VR 2019
- Steiner et al - Deep learning lymph node histopathology 2018
- Esteva et al - Deep learning in healthcare…[incl. surg. robotics] 2019

Papers compare methods (empirical!)

Hidden slide

- Olson et al – Large benchmark suite for ML evaluation… 2017
- Zhang et al – …comparison of state-of-the-art classification… 2017
- Delgado et al – Do we need hundreds of classifiers… 2014
- Caruana et al - …comparison…supervised learning… 2006
- Caruana et al – …evaluation…supervised…in high dimensions… 2006
- Caruana et al - …comparison…supervised…different…metrics 2003
- Lim et al - …thirty-three old and new classification… 2000

Books in AI/ML

Hidden slide

- Topol – Deep medicine 2019 (just starting – no opinion yet)
- Bishop – Pattern recognition and ML 2011
- Frank & Witten – Data mining: practical ML tools and techniques 4th ed. 2016, 3rd ed. 2011
- Shawe-Taylor & Cristianini – Kernel Methods for Pattern Analysis 2004
- Nixon & Aguado – Feature extraction and image processing 3rd ed. 2012
- Duda & Hart – Pattern Classification 2nd ed. 2000, 1st ed. 1973
- Russell & Norvig – Artificial intelligence 4th ed. 2020, 3rd ed. 2009

Books in biostatistics

Hidden slide

- Steyerberg – Clinical prediction models
- Harrell – Regression modeling strategies 2nd ed. , 1st ed.
- Urdan – Statistics in plain english 4th ed. 2016, 3rd ed. 2010
- Watkins – Introduction to the science of statistics 2016
- Zhou et al – Statistical methods in diagnostic medicine
- Hastie et al – Elements of statistical learning

Key authors of AI papers

AI

- Vladmir Vapnik
- Bernard Schölkopf
- Leo Brieman
- Isabelle Guyon
- Alexander Smola
- Robert Müller
- Corinna Cortes

Deep learning

- Geoffrey Hinton
- Yoshio Bengio
- Yan LeCun
- Andrew Wilson
- Ruslan Salakhutdinov

Key authors of statistical learning pa

Hidden slide

Statistics

- David Cox
- Donald Ruben
- Emanuel Parzen
- Frank Harrell Jr.
- Douglas Altman
- Leo Brieman

Biostatistics

- Frank Harrell Jr.
- Ewout Steyerberg
- Andrew Vickers
- Lisa Ohno-Machado
- Timothy Hastie
- Robert Tibshirani
- Douglas Manuel
- Peter Austin

Conferences

- Neural information processing society (NeurIPS/NIPS), ESANN
- International conference on ML (ICML), ECML, ACML
- Machine learning for health care
- Medical imaging consortium conference on AI (MICCAI)
- Knowledge discovery and data mining (KDD)
- Association of AI (AAAI)
- Uncertainty in AI (UAI)
- International joint conference on AI (IJCAI), ICMLA, IJCNN
- SIAM data mining (SDM)

Questions?

André Carrington PhD PEng

The Ottawa Hospital RESEARCH INSTITUTE | L'Hôpital d'Ottawa INSTITUT DE RECHERCHE

Affiliated with • Affilié à uOttawa